

Activelets and sparsity: A new way to detect brain activation from fMRI data

Ildar Khalidov¹, Dimitri Van De Ville¹, Jalal Fadili², and Michael Unser¹

¹Biomedical Imaging Group, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
²ENSI Caen, France

ABSTRACT

fMRI time course processing is traditionally performed using linear regression followed by statistical hypothesis testing. While this analysis method is robust against noise, it relies strongly on the signal model. In this paper, we propose a non-parametric framework that is based on two main ideas. First, we introduce a problem-specific type of wavelet basis, for which we coin the term “activelets”. The design of these wavelets is inspired by the form of the canonical hemodynamic response function. Second, we take advantage of sparsity-pursuing search techniques to find the most compact representation for the BOLD signal under investigation. The non-linear optimization allows to overcome the sensitivity-specificity trade-off that limits most standard techniques. Remarkably, the activelet framework does not require the knowledge of stimulus onset times; this property can be exploited to answer to new questions in neuroscience.

Keywords: fMRI, wavelets, exponential-spline wavelets, sparse approximations

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) has become a predominant modality for imaging human brain function. The blood-oxygenation-level-dependent (BOLD) signal, which originates from an overcompensation in oxygenated hemoglobine following neuronal activity by the neurovascular system, can be measured and localized using fast acquisition of T2* MRI volumes. The analysis of these large spatio-temporal datasets is difficult due to noise, measurement artifacts, and the variability of the BOLD response.

The classical approach puts forward a linear model to explain a voxel’s time course. The modeled BOLD response is the convolution of the hemodynamic response function (HRF) with the stimulus function, which can be a train of boxcar functions (block-based paradigm) or Dirac delta-functions (event-related paradigm). A widely used model for the HRF is the sum of two gamma-functions. Finally, voxels are declared as activated based on the results of statistical hypothesis testing.¹

The variability of the BOLD response is particularly important in event-related fMRI (efMRI). It has been shown that (1) the inter-subject differences in BOLD response are stronger than spatial ones inside one subject; (2) the intra-subject variability is stronger in space than in time; and (3) task-related responses have more coherent shapes than other activations, such as cardiac and respiratory ones.²

Several extensions have been proposed to the classical approach to deal with variations in BOLD response. Friston et al. include the orthonormalized versions of first partial derivatives of the HRF model with respect to parameters in the regression procedure.³ Adding more regressors reduces the number of degrees of freedom and requires looking for an optimal trade-off between sensitivity (ability to identify the active voxels) and specificity (ability to identify non-active voxels). Friman et al. use canonical correlation analysis to find basis functions having this optimal trade-off.⁴ To get ultimate adaptivity, one could build up the model for the HRF directly from the data. The FIR basis approach constructs the HRF as an FIR filter using statistical significance as an optimization criterion. Finally, we mention PCA and ICA methods that consider the problem as blind deconvolution of multiple sources. This strategy provides a data-dependent decomposition based on an information-theoretic criterion.

Further author information: (ildar.khalidov, dimitri.vandeville, michael.unser)@epfl.ch jalal.fadili@greyc.ensicaen.fr

In this paper, we propose a new dictionary of building blocks (activelets) that approximate well a class of gamma-function HRF models, instead of having a single HRF model. Recently, a family of exponential-spline wavelets has been introduced.⁵ Their properties are essentially transpositions of polynomial-related properties of traditional wavelets to the exponential-polynomial (gamma functions) setting. This makes exponential-spline wavelets natural candidates that can be tuned for our purpose.

Another important ingredient of our approach is the usage of sparsity-driven algorithms to recover (in a data-dependent way) the underlying activations. To get the minimum-budget building atoms of the event-related activation response, our problem is formulated as finding the sparsest solution to an underdetermined linear system of equations, which is the solution with the smallest number of nonzero elements. Therefore, we would like to solve a sparsity-promoting optimization problem, where a quadratic data fidelity term is optimized under a sparsity-budget constraint as measured by the l_1 norm. Solving such problems is the area of an active research work.

We validated this approach as follows. An initial experiment consisted into approximating arbitrary shifts of a single gamma function. For the noiseless case, we found that the exponential-spline wavelet basis allowed a very good approximation with as few as 4 coefficients. Interestingly, traditional B-spline wavelets required more than 30 coefficients to obtain the same approximation quality. To further demonstrate the feasibility of our approach, we evaluated its performance using a stochastic signal simulator. Three synthetic activations were generated as HRFs with random onset and parameters (time-to-peak and decay) and corrupted with additive white noise. The sparsity algorithms nicely detected all the activations within the presence of noise, even at SNRs at 15dB. Note that the activations were detected without any prior knowledge on the stimulus function, as opposed to model-based approaches. In addition, our method appears to cope well with the variability of BOLD responses.

This paper is organized as follows. In Section 2, we use the gamma-function model for the fMRI time series to find an appropriate wavelet (activelet) basis that could be used for the analysis. The activelet dictionary is used to formulate the sparse optimisation problem in Section 3. We outline our algorithm in Section 4 and show the validation results in Section 5. Finally, in Section 6, we discuss current challenges and the potential usage of the activelet method in neuroscience.

2. ACTIVELETS REPRESENTATION

Consider a number of neuronal events at the time points t_k , which can be modeled by a train of Dirac delta-functions as

$$s(t) = \sum_{k=1}^N \delta(t - t_k).$$

In fMRI, the timings t_k are known by the stimulation pattern. If we further assume the canonical HRF model with the linear-shift-invariant assumption, we can then expect the induced BOLD signal $x(t)$ to be

$$x(t) = s * h(t), \tag{1}$$

where $h(t)$ is the canonical HRF that is described by a difference of gamma functions:

$$h(t) = g(t; \tau_1, l_1) - \frac{1}{6}g(t; \tau_2, l_2), \quad g(t; \tau_l) = \frac{l^\tau t_+^{\tau-1} e^{-lt}}{\Gamma(\tau)}. \tag{2}$$

The parameters of the Gamma functions are chosen by specifying the delay to peak (τ/l) and the dispersion (τ/l^2). The standard parameters are: delay=6s, dispersion=1s for the initial peak, and delay=16s, dispersion=1s for the undershoot.¹ Therefore, $x(t)$ can be rewritten as $\sum_{k=1}^N h(t - t_k)$. The measured (noisy) data corresponds to

$$y[k] = x(k) + n[k], \tag{3}$$

where $n[k]$ represents the additive gaussian noise.

In traditional fMRI analysis, the relevant question is whether the voxel is active or not with respect to the known stimulus function. This question is answered by measuring the strength of the signal (1) in the data $y[k]$. A regression is fitted to the data; the resulting signal weight is subject to statistical testing for each voxel, using its corresponding t -value. In the case of efMRI, this method can be very sensitive to HRF variations: inaccurate estimation of the parameters τ_1 and τ_2 could lead to a false negative response. At the same time, τ_1 and τ_2 are fixed for the whole volume. In order to account for the variability of $h(t)$, state-of-the-art data analysis tools, such as SPM, include the partial derivatives dx/dt and $dx/d\tau_1$ in the regression. This makes the model more flexible; note, however, that each added regressor increases the risk of fitting non-stimuli related components in the signal, potentially giving way to false positive responses. Selection of regressors always boils down to setting a trade-off between sensitivity and specificity.

In this paper, we choose a different approach, which does not rely on the prior knowledge of the timings t_k . We note that the gamma-functions in (2) are piecewise exponential polynomials. A convenient tool to compactly represent such signals is the exponential-spline wavelet (ESW) transform.⁵ This transform is governed by a vector parameter $\vec{\alpha}$; similar to traditional wavelets, it provides a stable representation of $L_2(\mathbb{R}) = \oplus_{i=-\infty}^{\infty} W_i$ in terms of orthogonal wavelet spaces $W_i = \text{span}\{\psi_i(t - 2^i k)\}_{k \in \mathbb{Z}}$. As compared to the traditional case, the wavelets become scale-dependent. Nevertheless, the ESW transform can still be performed by means of fast Mallat's filterbank algorithm with scale-dependent filters.

Qualitatively, the ESW transform shows the same properties regarding exponential polynomials as the traditional wavelet transform has for polynomial functions. In particular, for piecewise exponential polynomials, it has the energy-compactness property, leading to sparse representation in the ESW domain. All these favourable properties of exponential-spline wavelets make them a natural tool for fMRI time-course analysis.

One degree of freedom is the choice of the ESW parameter $\vec{\alpha}$. We have optimized $\vec{\alpha}$ by minimizing the reconstruction error of the canonical HRF from K wavelet coefficients averaged over an interval of values for τ_1 and τ_2 . With the optimal choice of $\alpha = \{-0.6, -0.6, -0.6, -0.6\}$, we are able to reconstruct the HRF from as little as $K = 4$ coefficients with an average SNR of 30dB. We coin the term *activelets* the corresponding HRF-adapted exponential-spline wavelets. In comparison, to obtain the same quality using regular B-spline wavelets, which corresponds to setting $\vec{\alpha} = \{0, 0, 0, 0\}$, we need about 30 coefficients.

The activelet transform gives us a representation that concentrates the energy of an fMRI time course on few coefficients. In the next section, we describe a way to detect BOLD activity in the noisy signal by means of a sparse optimization problem.

3. OPTIMIZATION PROBLEM

We observe the noisy measurements $y[k]$ of the signal $x(t)$ as in (3), and we would like to find the sparsest set of wavelet coefficients β_0 that fits $y[k]$ best after reconstruction. Mathematically, we can state this as a constrained optimisation problem

$$\min_{\beta} \|\beta\|_0 \quad \text{subject to} \quad \|y - \Phi\beta\|_2 \leq \epsilon, \quad (4)$$

where Φ is the inverse ESW transform matrix, ϵ is a constant of the same order as the noise norm, and the ℓ_0 quasi-norm of a vector is the number of its non-zero components. Unfortunately, the functional to be optimized here is not convex, which makes the problem (4) NP-hard, demanding exorbitant computational efforts in general. This is why authors in⁶ advocated the convexified ℓ_1 form:

$$\min_{\beta} \|\beta\|_1 \quad \text{subject to} \quad \|y - \Phi\beta\|_2 \leq \epsilon, \quad (5)$$

Problem (5) can be also stated in the augmented lagrangian form

$$\beta_0 = \arg \min_{\beta} \|y - \Phi\beta\|_2^2 + \lambda \|\beta\|_1, \quad (6)$$

where λ is a regularisation parameter. The best choice of λ is determined by the actual noise level, which is estimated using e.g. the MAD.⁷

Table 1. Mean and standard deviation of random parameters in the synthetic experiment.

Parameter	Mean	Standard deviation
a_1, a_2, a_3	1	0.2
τ_1	6	1.5
τ_2	16	1.5

Note that the problems (6) and (4) are equivalent under an appropriate correspondence of parameters; that is, for a given $0 \leq \lambda < +\infty$ there exists $\epsilon > 0$ such that the two problems share the same solution, and vice versa. Formulation (6) is a (perturbed) linear programming problem known as BPDN.⁷

It has been shown in [6, Theorem 3.1] that if the noiseless ideal signal x is sufficiently sparse with $\|\beta\|_0 \leq (\mu_{\Phi}^{-1} + 1)/4^*$, then β_0 is the unique maximally sparse representation of x , and stable recovery is possible by minimizing (5). As far as the solution to (6) is concerned, Tropp [8, Theorem 8 and Corollary 9] proved that if x is sufficiently sparse in Φ with $\|\beta\|_0 < \mu_{\Phi}^{-1}/2$, and if the regularization parameter λ is chosen higher than the noise level, then the support of the solution to (6) is either contained or identical to the support of β_0 . In summary, if the signal has a sufficiently sparse representation, then, by minimizing either (5) or (6), we obtain a solution that (i) is the unique sparsest solution, (ii) is recovered with the appropriate sparsity pattern, and (iii) is stable to noise.

The optimisation in (6) can be cast as a (perturbed) linear program, and solved using general purpose solvers such as simplex and interior point methods which are slow for large-scale problems. There has also been interest in the statistical community in fitting regression models while imposing ℓ_1 -norm constraints on the regression coefficients. In,⁹⁻¹¹ the Lasso, LARS and Homotopy methods were proposed to solve (6) in the overdetermined case for all $\lambda \in [0, \|\Phi^T y\|_{\infty}]$. In this paper, we propose an approximate path-following algorithm which identifies groups of atoms at each iteration using a stagewise iterative-thresholding (IT) variant of LARS, where the sequence $\{\lambda_k\}_{k \geq 0}$ is not data-adapted, but allowed to be strictly decreasing.¹² It can be shown that the computational complexity of IT-LARS is $O(KV)$. Here, K is the number of non-zero coefficients in the solution, and V is the cost of the ESW decomposition.

We now have all the tools needed for the activelet method, which is outlined in the next section.

4. THE ACTIVELET APPROACH

The algorithm takes the fMRI time course $y[k]$ as input, which is fed into the sparse-solution search algorithm that uses the activelet decomposition routine to perform fast dictionary search. We use the undecimated wavelet (activelet) transform in order to gain shift-invariance. In addition, large-scale B-splines are incorporated in the dictionary to fit the fMRI baseline signal.

In case when the the voxel's time course encodes brain activity, the activelet part of the representation becomes non-zero but still remains sparse. The full reconstruction contains BOLD signal both due to activation and baseline. Contrarily, for a pure noisy voxel, coefficients of the activelet decomposition are kept to zero by the sparsity-pursuing algorithm and the fit reduces essentially to the baseline component.

5. SYNTHETIC EXPERIMENT

As a proof of concept, we have validated our approach on a synthetic dataset with 3 stimuli occurring at non-regularly spaced time moments. The variability is of HRF has been modelled by randomizing the amplitudes a_1, a_2, a_3 and the canonical HRF parameters τ_1, τ_2 of the 3 responses. The means and standard deviations of these variables in our experiment are shown in Table 1. The data was corrupted by gaussian white noise with standard deviation $\sigma = 0.15$.

In Figure 1, we show a typical activelet fit to a noisy dataset. We observe good visual quality of the reconstructed signal and a significant gain (+6 dB on average) in the SNR, despite the strong noise level in the original data. Importantly, as seen in Figure 2, a simulated pure-noise time course does not have a sparse activelet representation, and the algorithm detects no activations. We repeat the same experiment using traditional B-spline wavelets (cf. Fig. 3 and Fig. 4). The important degradation of results stresses the importance of a




Figure 1. Activelet fit to a noisy fMRI time-course. Thin line: original signal. Dashed line: noisy measurements. Thick line: reconstruction from sparse activelet representation.




Figure 2. Activelet fit (thick line) to pure noisy data (dashed line). No activations are detected.




Figure 3. Wavelet fit to a noisy fMRI time-course. Thin line: original signal. Dashed line: noisy measurements. Thick line: reconstruction from sparse wavelet representation.




Figure 4. B-spline wavelet fit (thick line) to pure noisy data (dashed line). Noise is partially fitted by the traditional wavelets.

properly designed, problem-adapted wavelet basis. In particular, sparse optimisation with traditional wavelets cannot distinguish pure noise from the activity-related signal.

6. CONCLUSION AND OUTLOOK

We have presented a novel approach for detecting activity in fMRI. The fMRI time course is decomposed in the newly designed activelet basis that concentrates activity-related energy on few wavelet coefficients; these coefficients are then identified by means of a sparse search algorithm.

The activelet representation and the sparse search algorithm are the two main ingredients for our method. From the sensitivity-specificity viewpoint, activelets alone are nothing but a very flexible model; used with a least-squares fit, they would “detect” the whole signal. At the same time, we have shown that sparse optimisation with traditional wavelets can not fit the signal properly.

Undecimated wavelet transform has a disadvantage of having highly correlated elements in the dictionary. This has a double negative impact on the algorithm: on one hand, the computation time amounts up to 5 sec/voxel; on the other hand, the noise could eventually well fitted by a linear combination of correlated elements. A potential improvement to the algorithm would consist in taking advantage of the knowledge of the dictionary structure and avoiding selection of correlated elements.

Interestingly, unlike traditional fMRI time-course processing methods, our algorithm does not use the knowledge of the stimulus times t_k . This opens a new horizon to the usage of fMRI in neuroscientific applications, such as mental chronometry, where t_k could be estimated from the time course itself.

ACKNOWLEDGMENTS

This work was supported by the Center for Biomedical Imaging (CIBM) of the Geneva - Lausanne Universities and the EPFL, the foundations Leenaards and Louis-Jeantet, as well as by the Swiss National Science Foundation under grant 200020-109415.

REFERENCES

1. R. Frackowiak, J. Ashburner, W. Penny, S. Zeki, and K. Friston, *Human Brain Function, Second Edition*, Academic Press, 2003.
2. M. Menz, J. Neumann, K. Mueller, and S. Zysset, “Variability of the BOLD response over time: An examination of within-session differences,” *NeuroImage* **32**, pp. 1185–1194, September 2006.
3. K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. Rugg, and R. Turner, “Event-related fMRI: Characterizing differential responses,” *NeuroImage* **7**, pp. 30–40, January 1998.
4. O. Friman, M. Borga, P. Lundberg, and H. Knutsson, “Adaptive analysis of fMRI data,” *NeuroImage* **19**(3), pp. 837–845, 2003.
5. I. Khalidov and M. Unser, “From differential equations to the construction of new wavelet-like bases,” *IEEE Trans. Signal Process.* **54**, pp. 1256–1267, April 2006.
6. D. Donoho, M. Elad, and V. Temlyakov, “Stable recovery of sparse overcomplete representations in the presence of noise,” *IEEE Trans. Inform. Th.* **52**, pp. 6–18, 2006.
7. S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM Journal on Scientific Computing* **20**(1), pp. 33–61, 1999.
8. J. Tropp, “Just relax: convex programming methods for identifying sparse signals in noise,” *IEEE Trans. Inform. Th.* **52**, pp. 1030–1051, March 2006.
9. R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society* **58**(1), pp. 267–288, 1996.
10. B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *Annals of Statistics* **32**(2), pp. 407–499, 2004.

* μ_Φ is the mutual coherence of Φ , that is the maximal off-diagonal element in magnitude of its Gram matrix, $1/\sqrt{T} \leq \mu_\Phi \leq 1$, where T is the number of data samples.

11. M. R. Osborne, B. Presnell, and B. A. Turlach, "A new approach to variable selection in least squares problems," *IMA J. Numerical Analysis* **20**, pp. 389–403, 2000.
12. M. J. Fadili and J.-L. Starck, "Sparse representation-based image deconvolution by iterative thresholding," in *Astronomical Data Analysis IV*, F. Murtagh and J.-L. Starck, eds., (Marseille, France), 2006.